LEGAL NOTICE:

This NVM Express revision 1.2 specification is proprietary to the NVM Express, Inc. (also referred to as "Company") and/or its successors and assigns.

NOTICE TO USERS WHO ARE NVM EXPRESS, INC. MEMBERS: Members of NVM Express, Inc. have the right to use and implement this NVM Express revision 1.2 specification subject, however, to the Member's continued compliance with the Company's Intellectual Property Policy and Bylaws and the Member's Participation Agreement.

NOTICE TO NON-MEMBERS OF NVM EXPRESS, INC.: If you are not a Member of NVM Express, Inc. and you have obtained a copy of this document, you only have a right to review this document or make reference to or cite this document. Any such references or citations to this document must acknowledge NVM Express, Inc. copyright ownership of this document. The proper copyright citation or reference is as follows: **"© 2007 - 2015 NVM Express, Inc. ALL RIGHTS RESERVED."** When making any such citations or references to this document you are not permitted to revise, alter, modify, make any derivatives of, or otherwise amend the referenced portion of this document in any way without the prior express written permission of NVM Express, Inc. Nothing contained in this document shall be deemed as granting you any kind of license to implement or use this document or the specification described therein, or any of its contents, either expressly or impliedly, or to any intellectual property owned or controlled by NVM Express, Inc., including, without limitation, any trademarks of NVM Express, Inc.

LEGAL DISCLAIMER:

THIS DOCUMENT AND THE INFORMATION CONTAINED HEREIN IS PROVIDED ON AN "**AS IS**" BASIS. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, NVM EXPRESS, INC. (ALONG WITH THE CONTRIBUTORS TO THIS DOCUMENT) HEREBY DISCLAIM ALL REPRESENTATIONS, WARRANTIES AND/OR COVENANTS, EITHER EXPRESS OR IMPLIED, STATUTORY OR AT COMMON LAW, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE, VALIDITY, AND/OR NONINFRINGEMENT.

All product names, trademarks, registered trademarks, and/or servicemarks may be claimed as the property of their respective owners.

NVM Express Workgroup
c/o Virtual, Inc.
401 Edgewater Place, Suite 600
Wakefield, MA 01880
info@nvmexpress.org

NVM Express™ Technical Errata

Errata ID	**003**
Revision Date	**6/3/2015**
Affected Spec Ver.	**NVM Express™ 1.2**
Corrected Spec Ver.	

Errata Author(s)

Name	Company
John Carroll	Intel
Peter Onufryk	PMC
Uma Parepalli	SKHMS
Judy Brock	Samsung
Ramdas Kachare	Seagate
Jason Gao	HGST

Errata Overview

The errata includes several editorial updates.

The errata clarifies protection information handling with metadata and end-to-end data protection.

The errata clarifies the RPMB feature, including MAC calculation.

The errata clarifies the under temperature threshold.

The errata clarifies the Host Memory Buffer feature, including the handshake for ownership of the buffer between the host and the controller.

Revision History

Revision Date	Change Description
1/29/2015	First draft.
2/26/2015	Added RPMB clarifications, and Figure 26 wording
4/2/2015	Filled in placeholders for changes and split portion of this into ECN 004
4/9/2015	Added missing authors and clarified default value for implementation specific. Edits during 4/9 call.
6/3/2015	Ratified.

Description of Specification Changes

Modify a portion of section 5.15 as shown below:

The settings specified in the Format NVM command are reported as part of the Identify Namespace data structure.

~~If the controller supports multiple namespaces, then t~~The host may specify the value of FFFFFFFFh for the namespace ID in order to apply the format operation to all namespaces accessible by the controller regardless of the value of the Format NVM Attribute field in the Identify Controller data structure.

The Format NVM command uses the Command Dword 10 field. All other command specific fields are reserved.

Modify a portion of section 6.6 as shown below:

The Compare command reads the logical blocks specified by the command from the medium and compares the data read to a comparison data buffer transferred as part of the command. If the data read from the controller and the comparison data buffer are equivalent with no miscompares, then the command completes successfully. If there is any miscompare, the command completes with an error of Compare Failure.

If metadata is provided, then a comparison is also performed for the metadata, excluding protection information. Refer to section 8.3.

Modify a portion of section 8.3 as shown below:

Figure TBD X illustrates the protection information processing that may occur as a side effect of Compare command processing. Compare command processing ~~is the same as that for a Read command except that no data is transferred to the host.~~ parallels both Write and Read commands. The controller checks the protection information contained in the command and the protection information read from the NVM.

Figure TBD X: Protection Information Processing for Compare

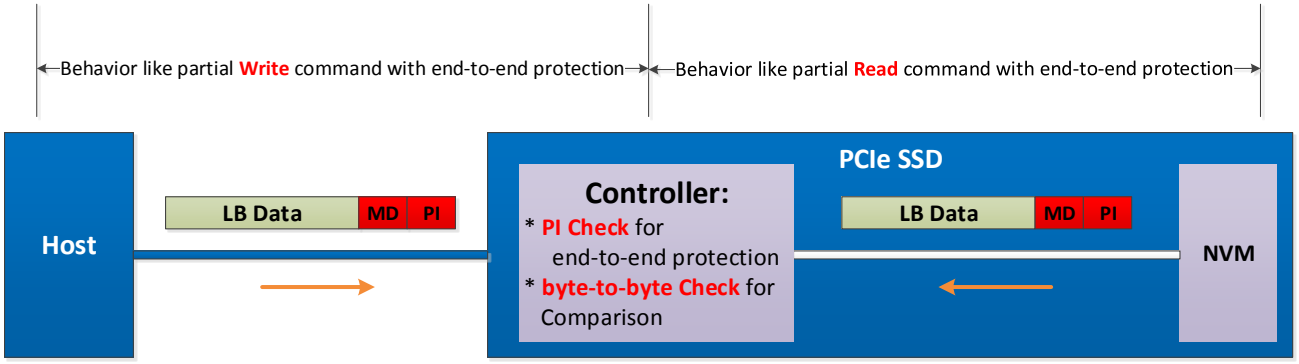

Protection Information with PRACT bit cleared to '0' (i.e., pass)

Modify a portion of section 8.3 as shown below:

The value of the computed reference tag for the first LBA of the command is the value contained in the Initial Logical Block Reference Tag (ILBRT) or Expected Initial Logical Block Reference Tag (EILBRT) field in the command, for writes and reads respectively. If the namespace is formatted for Type 1 or Type 2 protection, ~~T~~the computed reference tag is incremented for each subsequent logical block. If the namespace is formatted for Type 3 protection, the reference tag for each subsequent logic block remains the same as the initial reference tag. Unlike SCSI Protection Information Type 1 protection which implicitly uses the least significant four bytes of the LBA, ~~T~~the controller always uses the ILBRT or EILBRT field and requires host software to initialize the ILBRT or EILBRT field to the least significant four bytes of the LBA when Type 1 protection is used. In Type 1 protection, the controller should check the ILBRT or EILBRT field; if there is any miscompare, the command completes with an error of Invalid Protection Information.

Modify a portion of Figure 160 as shown below:

Figure 1: Compare – Command Dword 12

Bit	Description
29:26	**Protection Information Field (PRINFO):** Specifies the protection information action and check field, as defined in Figure 156. The Protection Information Action (PRACT) field shall be cleared to '0'.

Modify a portion of Figure 12 as shown below:

Figure 2: Command Format – NVM Command Set

<table>
<tr><th>Bytes</th><th>Description</th></tr>
<tr><td>39:24</td><td>
Data Pointer (DPTR): This field specifies the data used in the command.

If CDW0.PSDT ~~[15:14]~~ is set to 00b, then the definition of this field is:

<table>
<tr><td>39:32</td><td>PRP Entry 2 (PRP2): This field:
a) is reserved if the data transfer does not cross a memory page boundary.
b) specifies the Page Base Address of the second memory page if the data transfer crosses exactly one memory page boundary. E.g.,:
i. the command data transfer length is equal in size to one memory page and the offset portion of the PBAO field of PRP1 is non-zero or
ii. the Offset portion of the PBAO field of PRP1 is equal to zero and the command data transfer length is greater than one memory page and less than or equal to two memory pages in size.
c) is a PRP List pointer if the data transfer crosses more than one memory page boundary. E.g.,:
i. the command data transfer length is greater than or equal to two memory pages in size but the offset portion of the PBAO field of PRP1 is non-zero or
ii. the command data transfer length is equal in size to more than two memory pages and the Offset portion of the PBAO field of PRP1 is equal to zero.</td></tr>
<tr><td>31:24</td><td>PRP Entry 1 (PRP1): This field contains the first PRP entry for the command or a PRP List pointer depending on the command.</td></tr>
</table>

If CDW0.PSDT ~~[15:14]~~ is set to 01b or 10b, then the definition of this field is:

<table>
<tr><td>39:24</td><td>SGL Entry 1 (SGL1): This field contains the first SGL segment for the command. If the SGL segment is a Data Block descriptor, then it describes the entire data transfer. If more than one SGL segment is needed to describe the data transfer, then the first SGL segment is a Segment, or Last Segment descriptor. Refer to section 4.4 for the definition of SGL segments and descriptor types.</td></tr>
</table>
</td></tr>
<tr><td>23:16</td><td>
Metadata Pointer (MPTR): This field is valid only if the command has metadata that is not interleaved with the logical block data, as specified in the Format NVM command.

If CDW0.PSDT ~~[15:14]~~ is set to 00b, then this field shall contain the address of a contiguous physical buffer of metadata and shall be Dword aligned.

If CDW0.PSDT ~~[15:14]~~ is set to 01b, then this field shall contain the address of a contiguous physical buffer of metadata and shall be byte aligned.

If CDW0.PSDT ~~[15:14]~~ is set to 10b, then this field shall contain the address of an SGL segment containing exactly one SGL Descriptor and shall be Qword aligned. Refer to section 4.4.
</td></tr>
</table>

Modify a portion of section 2.1.2 as shown below:

02	RW	0	**Bus Master Enable (BME):** Enables the controller to act as a master for data transfers. When set to '1', bus master activity is allowed. When cleared to '0', the controller is not allowed to issue any Memory or I/O Requests~~, stops any active DMA engines, and returns to an idle condition~~.

Modify a portion of section 8.4 as shown below:

Associated with each power state is a Power State Descriptor in the Identify Controller data structure (refer to Figure 91). The descriptors for all implemented power states may be viewed as forming a table as shown in Figure 216 for a controller with seven implemented power states. Note that Figure 216 is illustrative and does not include all fields in the power state descriptor. The Maximum Power (MP) field indicates the instantaneous

maximum power that may be consumed in that state. The controller may employ autonomous power management techniques to reduce power consumption below this level, but under no circumstances is power allowed to exceed this level.

Figure 216: Example Power State Descriptor Table

Power State	Maximum Power (MP)	Entry Latency (ENTLAT)	Exit Latency (EXLAT)	Relative Read Throughput (RRT)	Relative Read Latency (RRL)	Relative Write Throughput (RWT)	Relative Write Latency (RWL)
0	25 W	5 µs	5 µs	0	0	0	0
1	18 W	5 µs	7 µs	0	0	1	0
2	18 W	5 µs	8 µs	1	0	0	0
3	15 W	20 µs	15 µs	2	0	2	0
4	10 W	20 µs	30 µs	1	1	3	0
5	8 W	50 µs	50 µs	2	2	4	0
6	5 W	20 µs	5000 µs	4	3	5	1

Modify a portion of Figure 26 as shown below:

Figure 3: Completion Queue Entry: DW 2

Bit	Description
31:16	**SQ Identifier (SQID):** Indicates the Submission Queue to which the associated command was issued ~~to~~. This field is used by host software when more than one Submission Queue shares a single Completion Queue to uniquely determine the command completed in combination with the Command Identifier (CID).

Modify a portion of section 8.10.2.3 as shown below:

The Authenticated Data Write is initiated by a Security Send command. The RPMB Data Frame delivered from the host to the controller includes the Request Message Type = 0003h, Block Count, Address, Write Counter, Data and MAC.

When the controller receives this RPMB Data Frame, it first checks whether the Write Counter has expired. If the Write Counter has expired then the controller sets the result to 0085h (write failure, write counter expired) and no data is written to the RPMB data area.

After checking the Write Counter is not expired, the Address is checked. If there is an error in the Address (e.g., out of range) then the result is set to 0004h (address failure) and no data is written to the RPMB data area.

After checking the Address is valid, the controller calculates the MAC (refer to section 8.10.1) ~~of Request Type, Block Count, Write Counter, Address and Data,~~ and compares this with the MAC in the request. If the MAC in the request and the calculated MAC are different, then the controller sets the result to 0002h (authentication failure) and no data is written to the RPMB data area.

If the MAC in the request and the calculated MAC are equal then the controller compares the Write Counter in the request with the Write Counter stored in the controller. If the counters are different then the controller sets the result to 03h (counter failure) and no data is written to the RPMB data area.

Modify a portion of section 8.10.2.4 as shown below:

The Authenticated Data Read sequence is initiated by a Security Send command. The RPMB data frame delivered from the host to the controller includes the Request Message Type = 0004h, Nonce, Address, and the Sector Count.

When the controller receives this RPMB Data Frame, it first checks the Address. If there is an error in the Address then the result is set to 0004h (address failure) and the data read is not valid.

~~After checking the Address is valid and successfully transferring the data to the host, the MAC is calculated from Response Type, Nonce, Address, Data and Result fields. If the MAC calculation fails then the returned result is 0002h (authentication failure).~~

When the host receives a successful completion of the Security Send command from the controller, it should send a Security Receive command to the controller to retrieve the data. The controller returns an RPMB Data Frame with Response Message Type (0400h), the Sector Count, a copy of the Nonce received in the request, the Address, the Data, the controller calculated MAC, and the Result. Note: It is the responsibility of the host to verify the MAC returned on an Authenticated Data Read Request.

If the data transfer from the addressed location in the controller fails, the returned Result is 0006h (read failure). If the Address provided in the Security Send command is not valid, then the returned Result is 0004h (address failure). If another error occurs during the read procedure then the returned Result is 0001h (general failure).

Modify a portion of Figure 224 as shown below:

Figure 224: RPMB Data Frame

Bytes	Component Name	Description
222-*N*:00	Stuff Bytes	Padding for the frame. Values in this field are not part of the MAC calculation. The size is ~~222~~ 223 bytes minus the size of the Authentication

Modify a portion of section 8.1as shown below:

1. The host issues a Firmware Image Download command to download the firmware image to the controller. There may be multiple portions of the firmware image to download, thus the offset for each portion of the firmware image being downloaded is specified in the Firmware Image Download command.
2. The host submits a Firmware Commit command with a Commit Action of 011b which specifies that the image should be activated immediately without reset. The downloaded image should replace the image in the firmware slot. If no image was downloaded since the last reset or Firmware Commit command, (i.e., the first step was skipped), then the controller shall verify and activate the image in the specified slot. If the controller starts to activate the firmware and Firmware Activation Notices are enabled (refer to Figure 122), the controller sends a Firmware Activation Starting asynchronous event to the host.

Modify a portion of section 5.13 as shown below:

The data structure used for the create operation is defined in Figure 102 and has the same format as the Identify Namespace data structure defined in Figure 92. After successful completion of a Namespace Management command with the create operation, the namespace is formatted with the specified attributes.

The fields that host software may specify in the create operation is defined in Figure 98. Fields that are reserved shall be cleared to 0h by host software. There is no data structure transferred for the delete operation.

Modify a portion of section 5.14.1.4 as shown below:

The default value of the over temperature threshold feature for Composite Temperature is the value in the Warning Composite Temperature Threshold (WCTEMP) field in the Identify Controller data if WCTEMP is non-zero; otherwise, it is implementation specific. The default value of the under temperature threshold feature for Composite Temperature is implementation specific. The default value of the over temperature threshold for all implemented temperature sensors is FFFFh. The default value of the under temperature threshold for all implemented ~~under~~ temperature sensors ~~thresholds~~ is 0h.

Modify a portion of section 5.14.1.13 as shown below:

This Feature controls the Host Memory Buffer. The attributes are indicated in Command Dword 11, Command Dword 12, Command Dword 13, Command Dword 14, and Command Dword 15.

The Host Memory Buffer feature provides a mechanism for the host to allocate a portion of host memory for the controller to use exclusively. After a successful completion of a Set Features enabling the host memory buffer, the host shall not write to the associated host memory region, buffer size, or descriptor list until the host memory buffer has been disabled.

<line break>

After a successful completion of a Set Features command that disables the host memory buffer, the controller shall not access any data in the host memory buffer until the host memory buffer has been enabled. The controller should retrieve any necessary data from the host memory buffer in use before posting the completion queue entry for the Set Feature command. Posting of the completion queue entry for the Set Feature command acknowledges that it is safe for the host software to modify the host memory buffer contents. Refer to section 8.9.

Modify a portion of section 8.9 as shown below:

The Host Memory Buffer feature allows the controller to utilize an assigned portion of host memory exclusively. The use of the host memory resources is vendor specific. Host software may not be able to provide any or a limited amount of the host memory resources requested by the controller. The controller shall function properly without host memory resources. Refer to section 5.14.1.13.

During initialization, host software may provide a descriptor list that describes a set of host memory address ranges for exclusive use by the controller. The host memory resources assigned are for the exclusive use of the controller (host software should not modify the ranges) until host software requests that the controller release the ranges and the controller completes the Set Features command. The controller is responsible for initializing the host memory resources. Host software should request that the controller release the assigned ranges prior to a shutdown event, a Runtime D3 event, or any other event that requires host software to reclaim the assigned ranges. After the controller acknowledges that it is no longer using the ranges, h~~H~~ost software may reclaim the host memory resources ~~after the controller acknowledges that it is no longer using the ranges~~. In the case of Runtime D3, host software should provide the host memory resources to the controller again and inform the controller that the ranges were in use prior to the RTD3 event and have not been modified.

The host memory resources are not persistent in the controller across a reset event. Host software should provide the previously allocated host memory resources to the controller after the reset completes. If host software is providing previously allocated host memory resources (with the same contents) to the controller, the Memory Return bit is set to '1' in the Set Features command.

The controller shall ensure that there is no data loss or data corruption in the event of a surprise removal while the Host Memory Buffer feature is being utilized.